

CLAIMS

WHAT IS CLAIMED IS:

1. A method for replicating data of a primary database system, comprising the steps of:
maintaining a buffer of transactions to be sent to a standby database system; and
synchronizing a transaction performed on the primary database system based on a number of transactions in the buffer and a predetermined number of transactions.
2. A method according to claim 1, wherein the step of synchronizing includes the step of:
blocking a commit of the transaction until the number of transactions in the buffers is in a
predetermined numerical relationship with the predetermined number of transactions.
3. A method according to claim 2, wherein:
the predetermined numerical relationship is less than.
4. A method according to claim 1, further comprising the step of:
executing a log writer process to record the transaction in a redo log.
5. A method according to claim 4, wherein:
the log writer process performs the step of synchronizing.
6. A method according to claim 4, wherein:
a database application process performs the step of synchronizing before submitting the
transaction to the log writer process.

7. A method according to claim 1, further comprising the step of:
executing a net server process to:
transmit the transaction over a network connection to the standby database system,
receive an acknowledgment that a redo record for the transaction has been written to a
standby log at the standby database system, and
remove the transaction from the buffer in response to the acknowledgment.
8. A method according to claim 1, further comprising the steps of:
receiving input from an operator indicating a transaction loss bound; and
setting the predetermined number of transactions based on the transaction loss bound.
9. A method according to claim 1, wherein the step of synchronizing includes the steps:
storing a counter indicating a number of the transactions in the buffer;
when adding the transaction to the buffer, incrementing the counter;
when removing the transaction from the buffer, decrementing the counter;
blocking a commit of the transaction when the counter is not less than the predetermined
number of transactions; and
acknowledging the commit of the transaction when the counter is less than the predetermined
number of transactions.
10. A computer-readable medium bearing instructions for causing one or more processors
to perform the steps of the method according to claim 1.
11. A method for replicating data of a primary database system, comprising the steps of:
maintaining a queue of transactions to be sent to a standby database system;
storing a counter indicating a number of the transactions in the queue;
storing a predetermined bound of transactions;
executing a log writer process to:

record the transaction in a redo log,
compare the counter and the predetermined bound,
if the counter is not less than the predetermined bound, then block a commit of the
transaction until the counter is less than the predetermined bound, and
if the counter is less than the predetermined bound, then increment the counter and
acknowledge the commit of the transaction; and
executing a net server process to:

transmit the transaction over a network connection to the standby database system,
receive an acknowledgment that a redo record for the transaction has been written to a
standby log at the standby database system, and
in response to the acknowledgment, remove the transaction from the queue and
decrement the counter.

12. A method for operating a log writer process to replicate data of a primary database system, comprising the steps of:

recording a transaction in a redo log;
comparing a counter indicating a number of the transactions in a queue of transactions to be
sent to a standby database system and a predetermined bound of transactions;
if the counter is not less than the predetermined bound, then blocking a commit of the
transaction until the counter is less than the predetermined bound, and
if the counter is less than the predetermined bound, then incrementing the counter and
acknowledging the commit of the transaction.

13. A computer-readable medium bearing instructions for causing one or more processors to perform the steps of the method according to claim 12.

14. A method for operating a net server process to replicate data of a primary database system, comprising the steps of:

accessing a transaction maintained in a buffer of transactions to be sent to a standby database system;
transmitting the transaction over a network connection to the standby database system;
receiving an acknowledgment that a redo record for the transaction has been written to a standby log at the standby database system, and
in response to the acknowledgment, removing the transaction from the queue and decrementing the counter.

15. A computer-readable medium bearing instructions for causing one or more processors to perform the steps of the method according to claim 14.

16. A method for replicating data in a primary database system having multiple database servers operating in parallel and accessing a common database on a shared disk, said method comprising the steps of:

setting a bound for each of the multiple database servers;
for each of the multiple database servers, performing the steps of:
maintaining a buffer of transactions to be sent to a standby database system; and
synchronizing a transaction performed on the primary database system based on a number of transactions in the buffer and the corresponding bound.